Supplement to "Attaining efficiency with imperfect public monitoring and one-sided Markov adverse selection"

(*Theoretical Economics*, Vol. 12, No. 3, September 2017, 957–978)

Daniel Barron
Kellogg School of Management, Northwestern University

Appendices A, B, and C contain all proofs for the results in Sections 3, 4, and 5, respectively. Appendix D contains a description of a plausible alternative equilibrium construction and a discussion of why this alternative cannot be used to prove Theorems 1 and 2.

A. Proofs for Section 3

Proof of Proposition 1. Many parts of this proof follow similar arguments in RSV. Suppose $\alpha \in \mathcal{A}^m \cup \mathcal{A}^M$ is implemented in an unrestricted mechanism. Then player 1 optimally tells the truth in each period and the result follows immediately.

Suppose $\alpha \in \mathcal{A}^I \setminus \{\mathcal{A}^m \cup \mathcal{A}^M\}$ is implemented in a T-period quota mechanism. For player $i \neq 1$,

$$\lim_{\delta \to 1} v_i^{(\delta,T)}(\alpha) = \lim_{\delta \to 1} E_{\theta \sim \pi}\left[\frac{1-\delta}{1-\delta^T}\sum_{t=0}^{T-1}\delta^t u_{i,t}\middle|\sigma\right] = \sum_{\theta \in \Theta}\frac{Q(\theta)}{T} g_i\big(\alpha(\theta),\theta\big)$$

because g_i is constant in θ. Moreover, $\lim_{T\to\infty}\frac{Q(\theta)}{T} = \pi(\theta)$, which proves the claim for $i \neq 1$.

Only player 1 takes actions, so her payoff is continuous in δ and $v_1^T(\alpha)$ is well defined. Following RSV, define the set of *copulas* $\mathcal{M} \subseteq \Delta(\Theta \times \Theta)$ as the set of distributions $\mu(m,\theta) \in \mathcal{M}$ such that $\sum_{m\in\Theta}\mu(m,\theta) = \sum_{m\in\Theta}\mu(\theta,m) = \pi(\theta)$ for all $\theta \in \Theta$. Let $\hat{\mu}_t(m,\theta) = \text{Prob}(m_t = m, \theta_t = \theta|\sigma)$ for some strategy σ.

I first claim that for any $\chi_1 > 0$, there exists some $T^* < \infty$ such that for any $T \geq T^*$, there exists a $\mu_\sigma \in \mathcal{M}$ such that

$$\left\|\frac{1}{T}\sum_{t=0}^{T-1}\hat{\mu}_t - \mu_\sigma\right\| < \chi_1. \tag{S9}$$

Note that $\sum_{m\in M}\sum_{t=0}^{T-1}\text{Prob}(m_t = m, \theta_t = \theta|\sigma) = \sum_{t=0}^{T-1}\text{Prob}(\theta_t = \theta|\sigma) = \sum_{t=0}^{T-1}\pi_t$. Because π is the stationary distribution of $P(\theta_{t+1}|\theta_t)$, $\lim_{T\to\infty}\frac{1}{T}\sum_{t=0}^{T-1}\pi_t = \pi$. Moreover,

Daniel Barron: d-barron@kellogg.northwestern.edu

DOI: 10.3982/TE1934

$\sum_{\theta\in\Theta}\sum_{t=0}^{T-1}\text{Prob}(m_t=m,\theta_t=\theta|\sigma)=\sum_{t=0}^{T-1}\text{Prob}(m_t=m|\sigma)=Q(m)$. Since $\lim_{T\to\infty}\frac{1}{T}\times\sum_{t=0}^{T-1}Q(m)=\pi(m)$, the marginals of the distribution $\text{Prob}(m_t=m,\theta_t=\theta|\sigma)$ converge to π as $T\to\infty$. It follows that there exists some copula $\mu_\sigma\in\mathcal{M}$ satisfying (S9). Further, the rate of convergence for each marginal is independent of the strategy σ.

In the limit as $\delta\to 1$, player 1's utility can be written

$$\begin{aligned}\frac{1}{T}\sum_{t=0}^{T-1}\sum_{\theta\in\Theta}\sum_{m\in\Theta}g_1\big(\alpha(m),\theta\big)\hat{\mu}_t(m,\theta)&=\sum_{\theta\in\Theta}\sum_{m\in\Theta}g_1\big(\alpha(m),\theta\big)*\frac{1}{T}\sum_{t=0}^{T-1}\hat{\mu}_t(m,\theta)\\&\le\sum_{\theta\in\Theta}\sum_{m\in\Theta}g_1\big(\alpha(m),\theta\big)\mu(m,\theta)+\chi_1|\Theta|^2\end{aligned}$$

by (S9).

Define $\mu^{\text{truth}}\in\mathcal{M}$ to be the copula satisfying $\mu^{\text{truth}}(\theta,\theta)=\pi(\theta)$ for all $\theta\in\Theta$, and $\mu^{\text{truth}}(m,\theta)=0$ otherwise. Lemma 1 from RSV can be slightly modified to show that an allocation rule satisfies (1) if and only if for all $\mu\in\mathcal{M}$,

$$\sum_{\theta\in\Theta}\sum_{m\in\Theta}g_1\big(\alpha(m),\theta\big)\mu(m,\theta)\le\sum_{\theta\in\Theta}\sum_{m\in\Theta}g_1\big(\alpha(\theta),\theta\big)\mu^{\text{truth}}(m,\theta).$$

Plugging in $\mu=\mu_\sigma$ and using (S9) yields

$$\frac{1}{T}\sum_{t=0}^{T-1}\sum_{\theta\in\Theta}\sum_{m\in\Theta}g_1\big(\alpha(m),\theta\big)\hat{\mu}_t(m,\theta)\le\sum_{\theta\in\Theta}\sum_{m\in\Theta}g_1\big(\alpha(\theta),\theta\big)\mu^{\text{truth}}(m,\theta)+\chi_1|\Theta|^2.$$

A similar argument can be applied to μ^{truth} to yield

$$\sum_{\theta\in\Theta}\sum_{m\in\Theta}g_1\big(\alpha(\theta),\theta\big)\mu^{\text{truth}}(m,\theta)\le\frac{1}{T}\sum_{t=0}^{T-1}\sum_{\theta\in\Theta}g_1\big(\alpha(\theta),\theta\big)\pi_t(\theta)+\chi_1|\Theta|^2.$$

Let σ be an optimal strategy. Consider a strategy $\tilde{\sigma}$ that reports truthfully as long as that message is available, and otherwise reports deterministically among the remaining types. This strategy is feasible and as $T\to\infty$, it can be shown that $\frac{1}{T}\sum_{t=0}^{T-1}\hat{\mu}_t\to\mu^{\text{truth}}$. Then there exists a T^* such that if $T\ge T^*$,

$$E\left[\frac{1}{T}\sum_{t=0}^{T-1}u_{1,t}\middle|\sigma\right]\le E_{\theta\sim\pi}\big[g_1\big(\alpha(\theta),\theta\big)\big]+\chi_1|\Theta|^2. \tag{S10}$$

But $\lim_{\delta\to 1}E[\frac{1-\delta}{1-\delta^T}\sum_{t=0}^{T-1}\delta^t u_{1,t}|\sigma]=E[\frac{1}{T}\sum_{t=0}^{T-1}u_{1,t}|\sigma]$, bounding $v_1^T(\alpha)$ from above.

Player 1 plays optimally and $\tilde{\sigma}$ is always feasible. So there exists a T^* such that if $T\ge T^*$,

$$E\left[\frac{1}{T}\sum_{t=0}^{T-1}u_{1,t}\middle|\sigma\right]\ge E_{\theta\sim\pi}\big[g_1\big(\alpha(\theta),\theta\big)\big]-2\chi_1|\Theta|^2. \tag{S11}$$

□

Proof of Proposition 2. Fix a (d,T) perturbed game, $\alpha\in\mathcal{A}^I$, and discount factor δ.

Step 1. Suppose that $\alpha \in \mathcal{A}^I \setminus \{\mathcal{A}^m \cup \mathcal{A}^M\}$ and so is implemented using a quota mechanism. Given a (d, T)-perturbed game, let $\sigma^{(d,T)}$ be player 1's optimal reporting strategy. For player 1,

$$\sigma^{(d,T)} \in \arg\max_{\sigma \in \Sigma} E\left[\frac{1-\delta}{1-\delta^T}\sum_{t=0}^{T-1}\delta^t u_{1,t} + \frac{1}{T}\sum d_i(h_t, \theta_T)\middle|\sigma\right].$$

In any strategy σ, message $m \in \Theta$ is sent $Q(m)$ times. In particular,

$$\begin{aligned} &E\left[\frac{1-\delta}{1-\delta^T}\sum_{t=0}^{T-1}\delta^t u_{1,t} + \frac{1}{T}\sum d_i(h_t, \theta_T)\middle|\sigma\right] \\ &\quad = E\left[\frac{1-\delta}{1-\delta^T}\sum_{t=0}^{T-1}\delta^t u_{1,t}\middle|\sigma\right] + \sum_{\theta \in \Theta}\frac{Q(\theta)}{T}E_{y_t,\theta_T}\big[d_i(h_t, \theta_T)|m_t = \theta\big]. \end{aligned}$$

The second term in this expression is constant in σ. Hence, $\sigma^{(d,T)}$ is an optimal strategy in the unperturbed game: $\sigma^{(d,T)} = \sigma^*_\delta(\alpha)$. Convergence to $v^T(\alpha)$ follows immediately.

Now suppose $\alpha \in \mathcal{A}^m \cup \mathcal{A}^M$ is implemented using an unrestricted mechanism. Consider a strategy σ that induces the same joint distribution over $(\theta_t, a_t)_{t=0}^{T-1}$ as σ^{truth}. Then Definition 7 implies that for all $t \le T$, $E[d(h_t^P, \theta_T)|\sigma] = E[d(h_t^P, \theta_T)|\sigma^{\text{truth}}]$, so player 1 cannot profitably deviate to σ. In particular, if α min- or max-maxes player $i \neq 1$, then $\alpha(\theta) = \alpha(\theta')$ for all $\theta, \theta' \in \Theta$ and player 1 has no profitable distribution from σ^{truth}.

Suppose α min- or max-maxes player 1, and suppose σ and σ^{truth} induce different joint distributions over $(\theta_t, a_t)_{t=0}^{T-1}$. Fix a history and type θ_t such that σ and σ^{truth} lead to different actions in period t for types θ_t. By Assumption 3, player 1 loses no less than $\frac{1-\delta}{1-\delta^T}\delta^t L$ in this period. She gains no more than $\frac{1}{T}\max_{h_t, \tilde{h}_t}\{d(h_t, \theta_T) - d(\tilde{h}_t, \theta_T)\} \le \frac{d}{T}$ at the end of the game by misreporting her type in period t. Continuation play is independent of period t because the mechanism is unrestricted. Therefore, player 1 has no incentive to lie in any period if

$$d \le \frac{1-\delta}{1-\delta^T}\delta^t L T \equiv d(\delta, T).$$

Under this condition, $\sigma^{(d,T)} = \sigma^*_\delta(\alpha) = \sigma^{\text{truth}}$.

Step 2. If $d < d(\delta, T)$, then $\sigma^{(d,T)} = \sigma^*_\delta(\alpha)$. Player 1's optimal strategy is independent of the prior ν, so payoffs are continuous in ν. Since $\lim_{\delta \to 1} v^{(\delta,T)}(\alpha) = v^T(\alpha)$ by definition, (3) holds for $\|\nu - \pi\|$ small and $\delta < 1$ close to 1. □

Proof of Proposition 3. Define $\sigma^{(k,j)} = \sigma^*_\delta(\alpha^{(k,j)})$ and let $\Sigma^{(k,j)}$ be the set of feasible strategies in the T-period mechanism implementing $\alpha^{(k,j)}$. Let h^t be a history at the beginning of block (k, j).

By Definition 8, only a single period $t^{(k,j)}$ chosen uniformly at random from $T^{(k,j)}$ will affect continuation play. Define

$$d_1(m, \theta) \equiv \frac{1}{1-\delta^T}\sum_{\hat{t}=t+T}^{\infty}\delta^{\hat{t}-t}(1-\delta)E\big[u_{1,t}|\sigma^*, h^t, \theta_{t+T} = \theta, m_{t^{(k,j)}} = m\big]. \qquad \text{(S12)}$$

Given θ_t, the distribution of θ_T is independent of h^t or any actions taken in block (k, j). Thus, σ^* is optimal if for any (k, j),

$$\sigma^{(k,j)} \in \arg\max_{\sigma\in\Sigma^{(k,j)}} E\left[\sum_{t'\in T^{(k,j)}} \left(\frac{1-\delta}{1-\delta^T}\delta^{t'-t}u_{1,t} + \frac{1}{T}d_1(m_{t'}, \theta_T)\right)\middle|\sigma, h^t\right]. \tag{S13}$$

By Proposition 2, $\sigma^{(k,j)}$ is an optimal strategy in any (d, T)-perturbed game with $d < d(\delta, T)$. Hence, to show σ^* optimal, it suffices to show that for any $m, \hat{m}, \theta \in \Theta$,

$$\left|d_1(m, \theta) - d_1(\hat{m}, \theta)\right| < d(\delta, T). \tag{S14}$$

Expression (S12) may be written

$$\begin{aligned} d_1(m, \theta) = &\sum_{k'=1}^{k}\sum_{j'=j+1}^{\infty}\sum_{t'\in T^{(k',j')}} \delta^{t'-t}\frac{1-\delta}{1-\delta^T}E\left[u_{1,t}|\sigma^*, h^t, \theta_{t+T} = \theta, m_{t^{(k,j)}} = m\right] \\ &+ \sum_{k'=k+1}^{K}\sum_{j'=j}^{\infty}\sum_{t'\in T^{(k',j')}} \delta^{t'-t}\frac{1-\delta}{1-\delta^T}E\left[u_{1,t}|\sigma^*, h^t, \theta_{t+T} = \theta, m_{t^{(k,j)}} = m\right]. \end{aligned}$$

Property (iii) of Definition 8 implies that for all $k' \neq k$ and $j' \geq j$, $\alpha^{(k',j')}$ does not depend on $m_{t^{(k,j)}}$. By definition of σ^*, actions in block (k', j') depend only on $\alpha^{(k',j')}$. Hence, for $k' \neq k$ and $j' \geq j$, actions and payoffs in block (k', j') are independent of $m_{t^{(k,j)}}$. So (S14) may be simplified to

$$\begin{aligned} &\left|d_1(m, \theta) - d_1(\hat{m}, \theta)\right| \\ &= \left|\sum_{j'=j+1}^{\infty}\sum_{t'\in T^{(k,j')}} \delta^{t'-t}\frac{1-\delta}{1-\delta^T}\left(E\left[u_{1,t}|\sigma^*, h^t, \theta_{t+T} = \theta, m_{t^{(k,j)}} = m\right]\right.\right. \\ &\qquad \left.\left. - E\left[u_{1,t}|\sigma^*, h^t, \theta_{t+T} = \theta, m_{t^{(k,j)}} = \hat{m}\right]\right)\right|. \end{aligned} \tag{S15}$$

If $\alpha(m) = \alpha(\hat{m})$, then $|d_1(m, \theta) - d_1(\hat{m}, \theta)| = 0$ by property (iii) of Definition 8.

Suppose $\alpha(m) \neq \alpha(\hat{m})$. An irreducible and aperiodic Markov chain converges to the invariant distribution at an exponential rate. So for any $\kappa > 0$ and $\epsilon > 0$, there exists a $K^* < \infty$ such that for any $K \geq K^*$, $j' > j$, and any prior $\nu \in \Delta(\Theta)$,

$$\|\pi_{\min T^{(k,j')} - \max T^{(k,j)}} - \pi\| < \kappa\epsilon^{j'-j}.$$

Payoffs satisfy $|u_1| \leq 1$, so

$$\begin{aligned} &\left|\sum_{t'\in T^{(k,j')}} \delta^{t'-t}\frac{1-\delta}{1-\delta^T}\left(E\left[u_{1,t}|\sigma^*, h^t, \theta_T = \theta, m_{t^{(k,j)}} = m\right]\right)\right. \\ &\left. - \sum_{t'\in T^{(k,j')}} \delta^{t'-t}\frac{1-\delta}{1-\delta^T}\left(E\left[u_{1,t}|\sigma^*, h^t, \pi_{\min T^{(k,j')}} = \pi, m_{t^{(k,j)}} = m\right]\right)\right| < \kappa\epsilon^{j'-j}. \end{aligned}$$

Moreover, by Proposition 1, there exists $\delta^* < 1$ such that if $\delta \geq \delta^*$, then

$$\begin{aligned}\Bigg|\sum_{t'\in T^{(k,j')}} \delta^{t'-t}\frac{1-\delta}{1-\delta^T}\big(E\big[u_{1,t}|\sigma^*, h^t, \pi_{\min T^{(k,j')}} = \pi, m_{t^{(k,j)}} = m\big]\big)\\ - \delta^{KTj'+kT}E\big[v_1^T(\alpha)|\sigma^*, h^t, m_{t^{(k,j)}} = m\big]\Bigg| < \epsilon.\end{aligned}$$

Combining these facts, for $K \geq K^*$ and $\delta \geq \delta^*$, gives

$$\begin{aligned}\Bigg|\sum_{t'\in T^{(k,j')}} \delta^{t'-t}\frac{1-\delta}{1-\delta^T}\big(E\big[u_{1,t}|\sigma^*, h^t, \theta_T = \theta, m_{t^{(k,j)}} = m\big]\big)\\ - \delta^{KTj'+kT}E\big[v_1^T(\alpha)|\sigma^*, h^t, m_{t^{(k,j)}} = m\big]\Bigg| \leq \kappa_1\epsilon^{j'-j} + \epsilon.\end{aligned} \tag{S16}$$

Combining (S16) with property (iv) of Definition 8, for any $m, \hat{m} \in \Theta$, yields

$$\big|d_1(m,\theta) - d_1(\hat{m},\theta)\big| < \delta^{kT}\epsilon\sum_{j'=0}^{\infty}\delta^{KTj'}2\big(\epsilon^{j'}\kappa + 1\big).$$

For any $\bar{\delta} < 1$, if $\delta^{TK} < \bar{\delta}$, then

$$\sum_{j'=0}^{\infty}\delta^{KTj'}\big(\epsilon^{j'}\kappa + 1\big) = \frac{1}{1-\delta^{TK}\epsilon}\kappa + \frac{1}{1-\delta^{TK}} < \frac{1}{1-\delta^{TK}\epsilon}\kappa + \frac{1}{1-\bar{\delta}}.$$

By Proposition 2, $\sigma^{(k,j)}$ is an optimal strategy in block (k,j) as long as

$$\delta^{kT}\epsilon\left(\frac{1}{1-\delta^{TK}\epsilon}\kappa + \frac{1}{1-\bar{\delta}}\right) < d(\delta, T).$$

For any $\delta \in (0,1)$, $d(\delta,T) > \delta^T L$. So this inequality holds for $\epsilon > 0$ sufficiently small. So for any $\bar{\delta} < 1$, there exists a $\delta^* < 1$ and $K^* < \infty$ such that for $K \geq K^*$, $\delta \geq \delta^*$, and $\delta^{TK} < \bar{\delta}$, σ^* is an optimal equilibrium of the (T,K)-recurrent mechanism.

At history h^t in block (k,j) and $j' > j$, (S16) implies (4). The inequality $\epsilon > 0$ may be made arbitrarily small by choosing δ^*, K^*, and $\bar{\delta}$ appropriately, proving the claim. □

B. Proofs for Section 4

Proof of Proposition 4. If v is (T,ζ,W,δ)-decomposable for some $\zeta > 0$, then it is (T,ζ',W,δ)-decomposable for any $\zeta' < \zeta$.

Fix $\epsilon > 0$, let $W \subseteq \mathbb{R}^N$ be a closed, convex, and bounded set, and consider $w \in W$. I construct an equilibrium with payoff $v \in B(w,\epsilon)$. From Proposition 3, choose $\underline{\delta} < 1$ and $\bar{K} > 0$ such that W is $(T,\zeta,\underline{\delta}^{T\bar{K}})$-decomposable and (4) holds for the chosen $\epsilon > 0$. Define $\bar{\delta} = \underline{\delta}^{\frac{\bar{K}}{\bar{K}+1}}$ and $\bar{\zeta} = \min_{\delta\in[\underline{\delta},\bar{\delta}]}\zeta(\delta^{T\bar{K}})$. Since ζ is continuous and strictly positive, $\bar{\zeta} > 0$. Then W is $(T,\bar{\zeta},\delta^{T\bar{K}})$-decomposable for any $\delta \in [\underline{\delta},\bar{\delta}]$. For any $\delta \geq \bar{\delta}$, there exists some $K \geq \bar{K}$ such that $\delta^{TK} \in [\underline{\delta}^{T\bar{K}}, \bar{\delta}^{T\bar{K}}]$ and (4) holds for a (T, K) mechanism and the given ϵ. So W is $(T, \bar{\zeta}, \delta^{TK})$-decomposable.

Given δ and K, define $\hat{\delta} = \delta^{TK}$. I construct a (T, K)-recurrent mechanism with payoffs that approximate w. For any $k \in \{0, \dots, K-1\}$, denote $\bar{w}^{(k,0)} \equiv w$, and let $\alpha^{(k,0)}$ and $w^{(k,0)}(y, \theta)$ be the allocation rule and the continuation payoff that $(T, \zeta, \hat{\delta})$-decompose $\bar{w}^{(k,0)}$.

Step 1. In block (k, j), implement an allocation rule $\alpha^{(k,j)}$ that $(T, \zeta, \hat{\delta})$-decomposes $\bar{w}^{(k,j)}$.[1] Let $w^{(k,j)}(y, \theta)$ be the corresponding continuation payoffs.

Step 2. At the end of block (k, j), choose one period $t^{(k,j)} \in T^{(k,j)}$ uniformly at random (using the public randomization device).

Step 3. In block $(k, j+1)$, set $\bar{w}^{(k,j+1)} \equiv w^{(k,j)}(y_{t^{(k,j)}}, \theta_{t^{(k,j)}})$.

I claim that the resulting mechanism is (T, K)-recurrent. Properties (i) and (ii) of Definition 8 follow immediately from the construction; property (iii) follows by noting that if $\alpha(\theta) = \alpha(\hat{\theta})$, then $w(y, \theta) = w(y, \hat{\theta})$ for all y. Thus, $w(y, \theta)$ and $w(y, \hat{\theta})$ can be implemented by the same sequence of allocation rules. For property (iv), given a public history h^t_{pub} at the beginning of block (k, j) and any $m_{t^{(k,j)}} \in \Theta$, (5) implies that $E[\bar{w}^{(k,j+1)}|\sigma^*, h^t_{\text{pub}}, m_{t^{(k,j)}}] = \bar{w}(h^t_{\text{pub}})$ is independent of $m_{t^{(k,j)}}$. Written nonrecursively,

$$\frac{1-\delta^T}{1-\delta^{TK}} \bar{w}(h^t_{\text{pub}}) = (1-\delta^T) \sum_{j'=1}^{\infty} (\delta^{TK})^{j'-1} E\big[v^T(\alpha^{(k,j+j')})|\sigma^*, h^t_{\text{pub}}, m_{t^{(k,j)}}\big]$$

for any $m_{t^{(k,j)}} \in \Theta$. So property (iv) of Definition 8 is also satisfied and this mechanism is (T, K)-recurrent.

If players commit to actions, then by Proposition 3, σ^* is an optimal strategy for player 1 and (4) holds. By definition,

$$\sum_{j=0}^{\infty} \hat{\delta}^{j'}(1-\hat{\delta}) E\big[v^T(\alpha^{(k,j)})|\sigma^*\big] \in B\big(w, (1-\hat{\delta}) v^T(\alpha^{(0,0)})\big)$$

since $v^T(\alpha^{(k,0)}) = v^T(\alpha^{(0,0)})$ for any $k \in \{0, \dots, K-1\}$. Because $v_i^T(\alpha^{(0,0)}) \in [-1, 1]$,

$$\sum_{j=0}^{\infty} \sum_{k=0}^{K-1} \delta^{KTj+kT} (1-\delta^T) E\big[v^T(\alpha^{(k,j)})|\sigma^*\big] \in B\big(w, (1-\hat{\delta})\big). \tag{S17}$$

Suppose $\underline{\delta}^{T\bar{K}}$ satisfies $1 - \underline{\delta}^{T\bar{K}} < \epsilon$. Then for each player i,

$$\begin{aligned}\sum_{t=0}^{\infty} \delta^t (1-\delta) E\big[u_{i,t}|\sigma^*\big] &= \sum_{j=0}^{\infty} \sum_{k=0}^{K-1} \sum_{t \in T^{(k,j)}} \delta^t (1-\delta) E\big[u_{i,t}|\sigma^*\big] \\ &\leq \sum_{j=0}^{\infty} \delta^{KTj} \sum_{k=0}^{K-1} (1-\delta^T) \delta^{kT} E\big[v^T(\alpha^{(k,j)}) + \epsilon|\sigma^*\big]\end{aligned}$$

[1] Using the public randomization device ξ to randomize among allocation rules, when appropriate.

$$\begin{aligned}&\leq w+(1-\hat{\delta})+\sum_{j=0}^{\infty}\delta^{KTj}\big(1-\delta^{KT}\big)\epsilon\\&=w+(1-\hat{\delta})+\epsilon\leq w+2\epsilon,\end{aligned}$$

where the first equality follows from rewriting the sum, the first inequality applies the upper bound from (4), the second inequality follows from (S17), and the final line follows immediately. A similar bound can be derived from below. Hence, the vector of equilibrium payoffs in the mechanism satisfies $v\in B(w,2\epsilon)$, as desired.

It remains to show that players have no incentive to deviate from the actions specified by the mechanism. Suppose that $\alpha^{(k,j)}\notin\mathcal{A}^M\cup\mathcal{A}^m$. A deviation at history h^t in block (k,j) affects payoffs in other blocks with probability $\frac{1}{T}$. Let $(1-\delta)B\in[0,1]$ be the maximum of 0 and the largest *myopic* gain for any type deviating from his equilibrium action.

By choice of K and δ, the gain from deviating to $a_t=a'$ at h^t is no more than

$$\begin{aligned}&(1-\delta)B+\frac{1}{T}\sum_{j'>j}\sum_{t'\in T^{(k,j')}}\delta^{t'}(1-\delta)\big(E\big[u_{1,t}|\sigma^*,h^t,t^{(k,j)}=t,a_t=a'\big]\\&\quad-E\big[u_{1,t}|\sigma^*,h^t,t^{(k,j)}=t,a_t=\alpha^{(k,j)}(m_t)\big]\big).\end{aligned}$$

Deviations are not profitable if this expression is weakly negative. Using (3), we can replace payoffs with invariant payoffs, plus an approximation error:

$$\begin{aligned}&(1-\delta)B+\frac{1}{T}\sum_{j'>j}\delta^{KT(j'-j)}\big(1-\delta^T\big)\big(\big(E\big[v_1^T\big(\alpha^{(k,j')}\big)|\sigma^*,h^t,t^{(k,j)}=t,a_t=a'\big]\\&\quad-E\big[v_1^T\big(\alpha^{(k,j')}\big)|\sigma^*,h^t,t^{(k,j)}=t,a_t=\alpha^{(k,j)}(m_t)\big]\big)+2\epsilon\big)\leq 0.\end{aligned}$$

Since $B\geq 0$, multiplying both sides by $\frac{1-\delta^{KT}}{1-\delta^T}$ and applying the definition of $\bar{w}$ and $w_i(y,m)$ yields the sufficient condition

$$\big(1-\delta^{TK}\big)B+\frac{\delta^{TK}}{T}\big(E_y\big[w_i(y,m)|a_i,\alpha_{-i}^{(k,j)}(m)\big]-\bar{w}_i\big)+2\frac{\delta^{TK}}{T}\epsilon\leq 0.$$

For $\epsilon>0$ sufficiently small, this inequality is implied by property (iii) of enforceability. Thus, for $\delta\geq\bar{\delta}$ and $K\geq\bar{K}$, player i has no profitable deviation from $\alpha^{(k,j)}\in\mathcal{A}^I\setminus\{\mathcal{A}^M\cup\mathcal{A}^m\}$.

If $\alpha^{(k,j)}\in\mathcal{A}^M\cup\mathcal{A}^m$ min- or max-maxes player $i\neq 1$, then the same argument proves the claim. If $\alpha^{(k,j)}$ min- or max-maxes player 1, then Proposition 3 proves that player 1 cannot profitably deviate by lying. I show that it is not profitable for player 1 either (a) to tell the truth and then deviate in action or (b) to misreport type and then deviate in action.

Consider the deviation (a). Property (iv) of enforceability implies that *conditional* on reporting truthfully, player 1 has no incentive to deviate in action.

Consider the deviation (b). Such a deviation does not change payoffs in subsequent periods of block (k, j), since the mechanism in that block is unconstrained. Therefore, the gain from such a deviation is bounded above by

$$\max_{m,a_1}\Big((1-\delta)\big(g_1\big(a_1, \alpha_{-1}(m), \theta\big) - g_1\big(\alpha(\theta), \theta\big)\big) + \frac{\delta^{TK}}{T}\big(E_y\big[w_1(y, m)|a_1, \alpha_{-1}(m)\big] - \bar{w}\big)\Big) + 2\frac{1}{T}\epsilon.$$

By definition of min-max and max-max, the first term in this expression is weakly negative. The value of the second term is independent of θ. Because (7) holds when $\theta = m$, $\frac{\delta^{TK}}{T}E_y[w_1(y, m)|a_1, \alpha_{-1}(m)] - \bar{w} \le -2\frac{1}{T}\epsilon$. So player 1 has no profitable deviation.

I have shown that for all $\delta \ge \underline{\delta}$, there exists a (T, K)-recurrent mechanism with payoffs satisfying $v \in B(w, \epsilon)$ that is also an equilibrium in the game without commitment. This proves the claim. □

Proof of Proposition 5. This proof proceeds in three steps: First, I define several key concepts; second, I prove a lemma that is a building block for the proposition; finally, I prove the proposition itself.

Definition S.1. For any $k \in \mathbb{R}$ and $\lambda \in \mathbb{R}^N$ such that $\|\lambda\| = 1$, define $H(\lambda, k) = \{v|\lambda \cdot v \le k\}$ as a *half-space in direction* λ. For fixed (T, ζ, δ), define the *maximal score attainable by allocation rule* $\alpha \in \mathcal{A}^I$ *in direction* λ, denoted $k^T(\alpha, \lambda, \zeta, \delta)$, as the maximum $k = \lambda \cdot v$ with $v \in \mathbb{R}^N$ such that v is $(T, \zeta, H(\lambda, k), \delta)$-decomposable with action α. Define $H^T(\alpha, \lambda, \zeta, \delta) = H(\lambda, k^T(\alpha, \lambda, \zeta, \delta))$.

Definition S.2. For a unit normal λ, define $k^T(\lambda) = \max_{\alpha\in\mathcal{A}^I} \lambda \cdot v^T(\alpha)$. Let $H^T(\lambda) = H(\lambda, k^T(\lambda))$ and $Q^T = \bigcap_\lambda H^T(\lambda)$.

Lemma S.1. *(i) There exists a continuous and decreasing function $\zeta(\delta)$ such that $k^T(\alpha, \lambda, \zeta(\delta), \delta)$ is independent of δ.*

(ii) Suppose λ is non-coordinate[2] *and let $\alpha \in \mathcal{A}^I$.Then for any $\zeta > 0$ and $\delta \in (0, 1)$, $k^T(\alpha, \lambda, \zeta, \delta) = \lambda \cdot v^T(\alpha)$.*

(iii) Let λ be coordinate with $\lambda_i = 1$ and $\alpha \in \mathcal{A}^m \cup \mathcal{A}^M$. Fix $\delta < 1$. For all $\epsilon > 0$, there exists a $\bar{\zeta} > 0$ such that if $\zeta \le \bar{\zeta}$, $k^T(\alpha, \lambda, \zeta, \delta) \ge \lambda \cdot v^T(\alpha) - \epsilon$.

Proof. (i) Fix a half-space H, and suppose $v \in H$ is (T, ζ, H, δ)-decomposable by allocation rule α and continuation payoffs $w(y, \theta)$. For any $\delta' \in (0, 1)$, define

$$w'(y, \theta) = \frac{\delta' - \delta}{\delta'(1-\delta)}v + \frac{\delta(1-\delta')}{\delta'(1-\delta)}w(y, \theta)$$

[2]A vector $\lambda \in \mathbb{R}^N$ is *coordinate* if exactly one element of λ is nonzero, and is otherwise *non-coordinate*.

and

$$\zeta' = \frac{1-\delta}{1-\delta'}\zeta.$$

The half-space H is convex and $v, w(y,\theta) \in H$, so $w'(y,\theta) \in H$. One can check that (6) and (7) hold for discount δ', so α and $w'(y,\theta)$ (T, ζ', H, δ')-decompose v.

(ii) This result is a natural modification of Lemma 5.4 in Fudenberg et al. (1994), with the sole difference that continuation payoffs satisfy (7).

(iii) Suppose λ is coordinate to the ith axis, so that for all $j \neq i$, $\lambda_j = 0$ and $\lambda_i = 1$. Let α min- or max-max player i. Fix $\delta < 1$. For $\epsilon > 0$, define $H^\epsilon = H(\lambda, \lambda \cdot v^T(\alpha) - \epsilon)$. I would like to show that for all $\epsilon > 0$, there exists a $\bar{\zeta} > 0$ such that if $\zeta \leq \bar{\zeta}$, then there exist $\{w(y,\theta)\}_{y,\theta} \subseteq H^\epsilon$ such that

$$\lambda \cdot \big((1-\delta)v^T(\alpha) + \delta E_{y,\theta}\big[w(y,\theta)|\alpha(\theta)\big]\big) \geq \lambda \cdot v^T(\alpha) - \epsilon, \tag{S18}$$

(7) holds for $j \neq i$, and for all $\theta \in \Theta$,

$$\frac{\delta}{T}\bar{w}_i - \zeta \geq \max_{a_i \in A_i \backslash \alpha_i(\theta)} \left\{ \frac{\delta}{T} E_y\big[w_i(y,\theta)|a_i, \alpha_{-i}(\theta)\big] \right\}.$$

For all $q \geq 0$, define the hyperplane $h^q = \{x \in \mathbb{R}^N | x_i = v_i^T(\alpha) - q\}$. For all $j \neq i$, let $\{w_j(y,\theta)\}_{y,\theta}$ satisfy (7) such that for all θ, $E_y[\hat{w}_j(y,\theta)|\alpha(\theta)] = v_j^T(\alpha)$. For any $\zeta > 0$, define $q^\zeta : Y \times \Theta \to \mathbb{R}$ such that $q^\zeta(y,\theta) \geq 0$ and for all θ, . $E_y[q^\zeta(y,\theta)|\alpha(\theta)] = \bar{q}^\zeta$ with

$$-\frac{\delta}{T}\bar{q}^\zeta - \zeta \geq -\max_{a_i \in A_i \backslash \alpha_i(\theta)} \left\{ \frac{\delta}{T} E_y\big[q^\zeta(y,\theta)\big]\Big| a_i, \alpha_{-i}(\theta) \right\}.$$

By pairwise full rank, such a q^ζ exists for any $\zeta > 0$. Moreover, $q^{\zeta'}(y,\theta) = \frac{\zeta'}{\zeta} q^\zeta(y,\theta)$. Let $w_i(y,\theta) = v_i^T(\alpha) - \epsilon - q^\zeta(y,\theta)$.

As defined above, $\{w(y,\theta)\}_{y,\theta} \in H^\epsilon$. Noting that $\lambda_j = 0$ and $\lambda_i = 1$, (S18) may be rewritten

$$(1-\delta)v_i^T(\alpha) + \delta E_{y,\theta}\big[v_i^T(a) - \epsilon - q^\zeta(y,\theta)|\alpha(\theta)\big] \geq v_i^T(\alpha) - \epsilon$$

or $E_{y,\theta}[q^\zeta(y,\theta)|\alpha(\theta)] \leq \frac{1-\delta}{\delta}\epsilon$. But $\lim_{\zeta \to 0} \max_{y,\theta} q^\zeta(y,\theta) = 0$ because $q^{\zeta'}(y,\theta) = \frac{\zeta'}{\zeta} q^\zeta(y,\theta)$. Therefore, for any $\epsilon > 0$, there exists a $\bar{\zeta} > 0$ such that this inequality holds if $\zeta \leq \bar{\zeta}$. This proves the claim. $\lhd$

Completing the Proof of Proposition 5. Let $W \subseteq Q^T$ be smooth. By Proposition 4, it suffices to show that for any $\delta \geq \hat{\delta}$, there exists a continuous function $\zeta(\delta) > 0$ such that W is $(T, \zeta(\delta), \delta)$-decomposable.

First, I claim that it suffices to show that for each $v \in W$, there exists a open set U_v with $v \in U_v$, a $\delta_v < 1$, and a continuous function $\zeta_v(\delta) > 0$ such that for any $\delta \geq \delta_v$, all $u \in U_v$ are $(T, \zeta_v(\delta), W, \delta)$-decomposable. Suppose for all $v \in W$ there exists such a δ_v and $\zeta_v(\delta)$. Then the set $\{U_v\}_{v \in W}$ is a open cover of W. Given that W is compact, there exists some finite subcover $\{U_{v_r}\}_{r=1}^R$. Let $\bar{\delta} = \max_r \delta_{v_r}$ and for all $\delta \geq \bar{\delta}$, let

$\zeta(\delta) = \min_r \zeta_{v_r}(\delta)$. Since $R < \infty$, $\bar{\delta} < 1$ and $\zeta(\delta) > 0$. Since each $\zeta_{v_r}(\delta)$ is continuous, $\zeta(\delta)$ is continuous. Then for all $\delta \geq \bar{\delta}$ and each $\zeta < \zeta(\delta)$, each U_{v_r} is $(T, \zeta(\delta), W, \delta)$-decomposable. Thus, $W \subseteq \bigcup_{r=1}^{R} U_{v_r}$ is $(T, \zeta(\delta), \delta)$-decomposable for continuous function $\zeta(\cdot)$.

Next, suppose that for each point on the boundary of W, $v \in \mathrm{bd}(W)$, there exists such an open set U_v with $v \in U_v$, $\delta_v < 1$, and continuous function $\zeta_v(\delta) > 0$. Then I claim that such U_v, $\delta_v < 1$, and $\zeta_v(\delta) > 0$ exist for every point $v \in W$. Because W is compact and convex, for all $v \in W$, there exist a finite number of points $\{v_1, \ldots, v_Z\} \in \mathrm{bd}(W)$ and weights $\{\gamma_1, \ldots, \gamma_Z\}$ such that $\sum \gamma_z = 1$ and $\sum \gamma_z v_z = v$. Take $\delta_v = \max_z \delta_{v_z}$, $\zeta_v(\delta) = \min_z \zeta_{v_z}(\delta)$, and $U_v = \{x | x = \sum_z \gamma_z x_z \text{ for } x_z \in U_{v_z}\}$. Then U_v is an open set with $v \in U_v$, and for $\delta \geq \delta_v$, U_v is $(T, \zeta_v(\delta), W, \delta)$-decomposable as a convex combination of decomposable points.

Finally, let $v \in \mathrm{bd}(W)$. I need to find an open set U_v with $v \in U_v$, a δ_v, and a continuous function $\zeta_v(\delta)$ such that U_v is $(T, \zeta_v(\delta), W, \delta)$-decomposable. Let λ be the unit normal to W at v, let $k = \lambda \cdot v$, and let $H = H^T(\lambda, k)$. Then $H^T(\lambda, k) \subseteq H^T(\lambda)$ holds strictly, since $W \subseteq \mathrm{int}(Q^T)$. Suppose λ is non-coordinate and let u be a boundary point of $H^T(\lambda)$. By Lemma S.1, for any $\zeta > 0$ and $\delta \in (0, 1)$, u is $(T, \zeta, H^T(\lambda), \delta)$-decomposable into allocation rule α and continuation payoffs $w(y, \theta) \in H^T(\lambda)$. Since v is a boundary point of H, which is a proper subset of $H^T(\lambda)$, there exists some $\delta < 1$, $\zeta > 0$, and $\epsilon > 0$ such that v can be $(T, \zeta, H^T(\lambda, k - \epsilon), \delta)$-decomposed using allocation rule α.

For $\delta' > \delta$, define $\zeta(\delta') = \frac{1-\delta'}{1-\delta}\zeta$. Note that $\zeta(\cdot)$ is continuous. Using (7), it can be shown that v can be $(T, \zeta(\delta'), H^T(\lambda, k - \frac{\delta(1-\delta')}{\delta'(1-\delta)}\epsilon), \delta')$-decomposed using allocation rule α. Moreover, the continuation payoffs $w'(y, \theta)$ satisfy $|w'(y, \theta) - v| \leq \bar{\kappa}(1 - \delta')$ for some $\bar{\kappa}$.

Define $U(\delta')$ as the ball around v of radius $2\bar{\kappa}(10\delta')$. Since W is smooth, for δ' sufficiently close to 1 there exists a $\tilde{\kappa} > 0$ such that the difference between H and W is at most $\tilde{\kappa}(1 - \delta')^2$. Hence, there exists some $\bar{\delta} < 1$ such that for all $\delta \geq \bar{\delta}$, there exists a $\zeta(\delta) > 0$ such that if $\zeta \leq \zeta(\delta)$, then v can be $(T, \zeta, \mathrm{int}(W), \delta)$-decomposed. Because continuation payoffs are in $\mathrm{int}(W)$, they can be translated by a constant independent of y to generate a neighborhood U_v about v that can be (T, ζ, W, δ)-decomposed.

Suppose now that λ is coordinate to the ith axis. As before, $H^T(\lambda, k) \subseteq H^T(\lambda)$ strictly. In particular, there exists some $\epsilon > 0$ such that $H^T(\lambda, k) \subseteq H^T(\lambda, v_i^T(\alpha) - \epsilon)$ strictly, where $\alpha \in \{\alpha^{m,i}, \alpha^{M,i}\}$. By Lemma S.1, for some $\delta < 1$ and $\zeta > 0$, a point on the boundary of $H^T(\lambda, v_i^T(\alpha) - \frac{\epsilon}{2})$ can be $(T, \zeta, H^T(\lambda, v_i^T(\alpha) - \frac{\epsilon}{2}), \delta)$-decomposed using allocation rule α. But then v can be $(T, \zeta, H^T(\lambda, v_i^T(\alpha) - \epsilon), \delta)$-decomposed using allocation rule a. The rest of the proof proceeds as in the previous case.

Fix $\alpha \in \mathcal{A}^I$ and suppose there exists some λ such that $v^T(\alpha) \notin H(\lambda)$. Then by definition, $\lambda \cdot v^T(\alpha) > \max_{\alpha' \in \mathcal{A}^I} \lambda \cdot v^T(\alpha')$, a contradiction. The set Q^T is convex, so $V^{T*} \subseteq Q^T$. Hence, any smooth $W \subseteq \mathrm{int}(V^{T*})$ can be approximated by a set of equilibrium payoffs for sufficiently high δ. $\square$

C. Proofs for Section 5

For the purposes of this proof, assume without loss that $E_{\theta\sim\pi}[\tau^{\text{truth}}(\theta)]=0$ and define $\hat{L}=\min\{\tilde{L},\bar{L}\}>0$.

Definition S.3. Let $\alpha\in\{\hat{\alpha}^{M,i},\hat{\alpha}^{m,i}\}_{i=2}^{N}$. Define the *invariant payoff* for α as

$$\hat{v}^T(\alpha)=E\left[\sum_{t=0}^{T-1}\delta^t(1-\delta)g\big(\alpha(\theta),\theta\big)\right].$$

For all other allocation rules $\alpha\in\hat{\mathcal{A}}^I$, let $\hat{v}^T(\alpha)=v^T(\alpha)$. Define $\hat{V}^{T*}$ analogously to V^{T*}:

$$\hat{V}^{T*}=\text{co}\big\{\hat{v}^T(\alpha)|\alpha\in\hat{\mathcal{A}}^I,\text{for all } i:\hat{v}_i^T(\alpha)\geq\hat{v}_i^T\big(\alpha^{m,i}\big)\big\}.$$

Lemma S.2. *Suppose either of the following statements:*

(i) The allocation rule $\alpha\in\hat{\mathcal{A}}^I\setminus\{\hat{\mathcal{A}}^m\cup\hat{\mathcal{A}}^M\}$ is implemented by a T-period quota mechanism.

(ii) The allocation rule $\alpha\in\{\hat{\alpha}^{M,1},\hat{\alpha}^{m,1}\}$ is implemented by a T-period unrestricted mechanism.

Define $\hat{d}(\delta,T)=2\frac{1-\delta}{1-\delta^T}\delta^T T\hat{L}$. There exists a $\bar{\delta}<1$ such that for any $\delta\geq\bar{\delta}$ and $d<\hat{d}(\delta,T)$, player 1's optimal strategy equals $\sigma_\delta^*(\alpha)$. Moreover, for all $\epsilon>0$, there exists $\chi>0$ and $\delta^*<1$ such that if $\|\nu-\pi\|<\chi$, $\delta>\delta^*$, and $d<\hat{d}(\delta,T)$, then

$$E\left[\frac{1-\delta}{1-\delta^T}\sum_{t=0}^{T-1}\delta^t u_t\middle|\sigma^{(d,T)}\right]\in B\big(\hat{v}^T(\alpha),\epsilon\big). \tag{S19}$$

Let σ^{truth} be the strategy in which $m_t=\theta_t$ for all $t\geq 0$. Then $\sigma_\delta^*(\alpha)=\sigma^{\text{truth}}$ for $\alpha\in\{\hat{\alpha}^{m,1},\hat{\alpha}^{M,1}\}$.

Proof. Suppose $\alpha\in\{\hat{\alpha}^{M,1},\hat{\alpha}^{m,1}\}$. As in Lemma S.2, $\sigma^{(d,T)}=\sigma^{\text{truth}}$ if $d<d(\delta,T)$. The inclusion (S19) follows immediately.

Suppose $\alpha\in\hat{\mathcal{A}}^I\setminus\{\hat{\mathcal{A}}^m\cup\hat{\mathcal{A}}^M\}$ is implemented by a T-period quota mechanism. By definition,

$$\sigma^{(d,T)}(\delta)\in\arg\max_{\sigma}E\left[\frac{1-\delta}{1-\delta^T}\sum_{t=0}^{T-1}\delta^t g\big(\alpha(\theta_t),\theta_t\big)+\frac{1}{T}\sum_{t=0}^{T-1}d\big(h_t^P,\theta_T\big)\middle|\sigma\right].$$

As in Lemma S.2, each $m\in\Theta$ is sent exactly $Q(m)$ times. Therefore, $\sigma_\delta^*(\alpha)$ maximizes player 1's payoff. The unperturbed game may be written as a Markov decision problem, so there exists a Blackwell optimal strategy $\sigma^*(\alpha)$ and a threshold $\delta^*<1$ such that for any $\delta>\delta^*$, $\sigma_\delta^*(\alpha)=\sigma^*(\alpha)$. It follows that

$$\lim_{\delta\to 1}E\left[\frac{1-\delta}{1-\delta^T}\sum_{t=0}^{T-1}\delta^t u_t\middle|\sigma^{(d,T)}\right]=v^T(\alpha).$$

□

Definition S.4. For $\tau : \Theta \to \mathbb{R}$, $d \in \mathbb{R}_+$, and $T \in \mathbb{N}$, a (d,T)*-perturbed game simulating transfers* τ is a T-period game with payoffs (2). For any $m_t, \hat{m}_t, \theta_T \in \Theta$,

$$E_{y_t}\big[d_1\big((m_t, y_t), \theta_T\big)|\alpha(m_t)\big] - E_{y_t}\big[d_1\big((\hat{m}_t, y_t), \theta_T\big)|\alpha(\hat{m}_t)\big] \in B\big(\tau(m_t) - \tau(\hat{m}_t), d\big).$$

For any $m_t = \hat{m}_t$ such that $\alpha(m) = \alpha(\hat{m})$ and $y = \hat{y}$, $d((m_t, y_t), \theta_T) = d((\hat{m}_t, y_t), \theta_T)$.

Lemma S.3. *Suppose* $\alpha \in \{\hat{\alpha}^{M,i}, \hat{\alpha}^{m,i}\}_{i=2}^N$. *Then player* 1*'s optimal reporting strategy equals* σ^{truth} *in any* (d,T)*-perturbed game simulating transfers* τ^{truth} *with* $d < \hat{d}(\delta, T)$. *For any* $\epsilon > 0$, *there exists a* $\bar{\delta} < 1$ *and* $\chi > 0$ *such that if* $\delta \geq \bar{\delta}$ *and* $\|\nu - \pi\| < \chi$,

$$E\left[\frac{1-\delta}{1-\delta^T}\sum_{t=0}^{T-1}\delta^t u_t \middle| \sigma^{(d,T)}\right] \in B\big(\hat{v}^T(\alpha), \epsilon\big).$$

Proof. By Assumption 4, for any $\theta, \theta' \in \Theta$ such that $\alpha(\theta) \neq \alpha(\theta')$, $g_1(\alpha(\theta), \theta) - \tau^{\text{truth}}(\theta) - \hat{L} > g_1(\alpha(\theta'), \theta) - \tau^{\text{truth}}(\theta')$. Consider a (d,T)-perturbed game with transfers τ^{truth} satisfying $d < \hat{d}(\delta, T)$, and let $\sigma^{(d,T)}$ be player 1's optimal reporting strategy. Continuation payoffs are independent of history in an unrestricted mechanism, so $\sigma^{(d,T)} = \sigma^{\text{truth}}$ if for any $t \leq T$, for all $m, m', \theta, \theta_T \in \Theta$,

$$\begin{aligned}&\frac{1-\delta}{1-\delta^T}\delta^t g_1\big(\alpha(m), \theta\big) + \frac{1}{T}E_y\big[d\big((m,y), \theta_T\big)|\alpha(m)\big]\\ &\quad\geq \frac{1-\delta}{1-\delta^T}\delta^t g_1\big(\alpha(m'), \theta\big) + \frac{1}{T}E_y\big[d\big((m',y), \theta_T\big)|\alpha(m')\big].\end{aligned}$$

This condition trivially holds if $\alpha(\theta) = \alpha(\theta')$. If $\alpha(\theta) \neq \alpha(\theta')$, then $E_y[d((m,y), \theta_T)|\alpha(m)] - E_{y_t}[d((m',y), \theta_T)|\alpha(m')] \geq \tau(\theta) - \tau(\theta') - 2\hat{d}$. So player 1 reports truthfully if

$$\frac{1-\delta}{1-\delta^T}\delta^t g_1\big(\alpha(\theta), \theta\big) - \frac{2\hat{d}}{T} + \frac{1}{T}\tau(\theta) \geq \frac{1-\delta}{1-\delta^T}\delta^t g_1\big(\alpha(\theta'), \theta\big) + \frac{1}{T}\tau\big(\theta'\big).$$

This inequality holds strictly as $\delta \to 1$ because $\lim_{\delta\to1}\frac{1-\delta}{1-\delta^T}\delta^t = \frac{1}{T}$. So there exists $\bar{\delta} < 1$ such that for all $\delta \geq \bar{\delta}$, player 1 reports truthfully in each period. Hence,

$$\lim_{\delta\to1} E\left[\frac{1-\delta}{1-\delta^T}\sum_{t=0}^{T-1}\delta^t u_t \middle| \sigma^{(d,T)}(\delta)\right] = v^T(\alpha)$$

as desired. □

Lemma S.4. *For any* $\alpha \in \hat{\mathcal{A}}^I$,

$$\lim_{T\to\infty}\hat{v}^T(\alpha) = E_{\theta\sim\pi}\big[g\big(\alpha(\theta), \theta\big)\big].$$

Proof. Suppose first that $\alpha \in \{\hat{\alpha}^{M,i}, \hat{\alpha}^{m,i}\}_{i=1}^N$. If $i = 1$, then $\sigma_\delta^*(\alpha) = \sigma^{\text{truth}}$ by Lemma S.2 and the result follows immediately. If $i \neq 1$, the result follows immediately by Definition S.3.

Suppose instead that $\alpha \in \hat{\mathcal{A}}^I \setminus \{\hat{\alpha}^{M,i}, \hat{\alpha}^{m,i}\}_{i=1}^N$. This argument borrows heavily from the analogous argument by RSV. I claim that player 1 reports truthfully "with high probability in each period." Formally, for any $\chi > 0$, there exists a $T^* < \infty$ such that if $T \geq T^*$, for all $\theta \in \Theta$,

$$\lim_{\delta \to 1} \frac{1}{T} \sum_{t=0}^{T-1} \operatorname{Prob}\{\alpha(m_t) = \alpha(\theta_t) | \sigma\} > 1 - \chi. \tag{S20}$$

Toward contradiction, suppose there exists $\chi > 0$ such that for all T^*, (S20) does not hold. Recall the set of *copulas* $\mathcal{M} \subseteq \Delta(\Theta \times \Theta)$ from Proposition 1. Define μ as the copula that player 1's optimal strategy σ approximates. For both (S10) and (S11) to hold simultaneously, it must be that

$$4\chi_1 |\Theta|^2 \geq \sum_{\theta \in \Theta} \sum_{m \in \Theta} g_1(\alpha(m), \theta)(\mu^{\text{truth}}(m, \theta) - \mu(m, \theta)). \tag{S21}$$

I seek to bound $\|\mu^{\text{truth}} - \mu\|$ using this statement.

Given $\alpha \in \hat{\mathcal{A}}^{\text{CM}}$, Lemma 1 from RSV can be slightly modified to show that

$$\sum_{\theta \in \Theta} \sum_{m \in \Theta} \mu(m, \theta) g_1(\alpha(m), \theta) < \sum_{\theta \in \Theta} \sum_{m \in \Theta} \mu^{\text{truth}}(m, \theta) g_1(\alpha(\theta), \theta)$$

for any μ that assigns positive weight to (m, θ) combinations for which $\alpha(m) \neq \alpha(\theta)$. The variable μ_0 is one of the finite number of extremal points $\{\mu_0, \ldots, \mu_R\} \subseteq \mathcal{M}$. For each $r \leq R$, either $\mu_r(m, \theta) > 0$ for $m, \theta \in \Theta$ only if $a(m) = a(\theta)$, or there exists $c > 0$ such that for any $m, \theta \in \Theta$,

$$\sum_{\theta \in \Theta} \sum_{m \in \Theta} \mu^{\text{truth}}(m, \theta) g_1(\alpha(\theta), \theta) - \sum_{\theta \in \Theta} \sum_{m \in \Theta} \mu_r(m, \theta) g_1(\alpha(m), \theta) > c.$$

Let $\mathcal{R}_T = \{r \leq R | \mu_r(m, \theta) > 0 \text{ only if } \alpha(m) = \alpha(\theta)\}$. There exist $\beta_r \geq 0$ that sum to 1 such that $\mu = \sum_r \beta_r \mu_r$. Therefore, (S21) may be written

$$\begin{aligned}
4\chi_1 |\Theta|^2 &\geq \sum_{\theta \in \Theta} \sum_{m \in \Theta} g_1(\alpha(m), \theta) \left(\mu^{\text{truth}}(m, \theta) - \sum_{r \leq R} \beta_r \mu_r(m, \theta) \right) \\
&= \sum_{\theta \in \Theta} \sum_{m \in \Theta} g_1(\alpha(m), \theta) \left((1 - \beta_0) \mu^{\text{truth}}(m, \theta) - \sum_{r=1}^{R} \beta_r \mu_r(m, \theta) \right) \\
&= \sum_{\theta \in \Theta} \sum_{m \in \Theta} g_1(\alpha(m), \theta) \sum_{r=1}^{R} \beta_r (\mu^{\text{truth}}(m, \theta) - \mu_r(m, \theta)) \\
&= \sum_{r \notin \mathcal{R}_T} \beta_r \left(\sum_{\theta \in \Theta} \sum_{m \in \Theta} (g_1(\alpha(m), \theta)(\mu^{\text{truth}}(m, \theta) - \mu_r(m, \theta))) \right) > c \sum_{r \notin \mathcal{R}_T} \beta_r.
\end{aligned}$$

So (S21) implies

$$4\chi_1 |\Theta|^2 > c \sum_{r \notin \mathcal{R}_T} \beta_r.$$

For any $\chi_2 > 0$, there exists $T^* < \infty$ such that for all $T \geq T^*$, there exists $\delta^* < 1$ such that for all $\delta \geq \delta^*$, $\chi_1 |\Theta|^2 < \chi_2$ and, thus, $\sum_{r \notin \mathcal{R}_T} \beta_r < \frac{\chi_2}{c}$.

For any $\mu : \Theta \times \Theta \to \mathbb{R}$, define $\lambda(\mu) = \sum_{(m,\theta)|a(m)=a(\theta)} \mu(m,\theta)$. Then $\lambda(\mu) \geq 1 - \sum_{r \notin \mathcal{R}_T} \beta_r > 1 - \frac{\chi_2}{c}$. By (S9), for any $(m,\theta) \in \Theta \times \Theta$, $\frac{1}{T} \sum_{t=0}^{T-1} \hat{\mu}_t(m,\theta) - \mu(m,\theta) < \chi_1$. Therefore, $|\lambda(\frac{1}{T} \sum_{t=0}^{T-1} \hat{\mu}_t) - \lambda(\mu)| < \chi_1 |\Theta|^2$ and so

$$\frac{1}{T} \sum_{t=0}^{T-1} \text{Prob}\big\{\alpha(m_t) = \alpha(\theta_t) | \sigma\big\} > 1 - \frac{\chi_2}{c} - \chi_1 |\Theta|^2.$$

As $\delta \to 1$, choosing $\chi_1, \chi_2 > 0$ so that $\frac{\chi_2}{c} + \chi_1 |\Theta|^2 < \chi$ proves the contradiction.

For any $\epsilon > 0$, choose $T^* < \infty$ such that for any $T \geq T^*$, there exists $\delta^* < 1$ such that if $\delta \geq \delta^*$,

$$\frac{1-\delta}{1-\delta^T} \sum_{t=0}^{T-1} \delta^t \, \text{Prob}\big\{\alpha(m_t) = \alpha(\theta_t) | \sigma^{(0,T)}\big\} > 1 - \epsilon.$$

The claim follows. □

Definition S.5. Consider the infinite-horizon dynamic game and fix $T, K \in \mathbb{N}$ and $\delta \in (0,1)$. A (T,K)*-recurrent mechanism in the game with an expert* satisfies the properties of Definition 8, with the following changes:

- The allocation implemented in block (k,j) is $\alpha^{(k,j)} \in \hat{\mathcal{A}}^I$.
- For any public history h^t_{pub} at the beginning of block (k,j), property (iv) is replaced by the following conditions:
  (i) If $\alpha^{(k,j)} \in \hat{\mathcal{A}}^I \setminus \{\hat{\alpha}^{m,i}, \hat{\alpha}^{M,i}\}_{i=2}^N$, then there exists $\bar{w}_1^{(k,j)}(h^t_{\text{pub}})$ such that for all $t^{(k,j)} \in T^{(k,j)}$, $m_{t^{(k,j)}} \in \Theta$,
  $$\sum_{j'=j+1}^{\infty} \delta^{TK(j'-j)} (1-\delta^{TK}) E\big[v_1^T\big(\alpha^{(k,j')}\big) | \sigma^*, h^t_{\text{pub}}, m_{t^{(k,j)}}, \pi_{\min T^{(k,j)}} = \pi\big]$$
  $$= \bar{w}_1^{(k,j)}\big(h^t_{\text{pub}}\big).$$
  (ii) If $\alpha^{(k,j)} \in \{\hat{\alpha}^{m,i}, \hat{\alpha}^{M,i}\}_{i=2}^N$, then define $\tau^{(k,j)} : \Theta \to \mathbb{R}$ as the transfers satisfying (4) for the allocation rule $\alpha^{(k,j)}$. Then there exists $\bar{w}_1^{(k,j)}(h^t_{\text{pub}}) \in \mathbb{R}$ such that for all $\theta \in \Theta$,
  $$\sum_{j'=j+1}^{\infty} \delta^{TK(j'-j)} (1-\delta^{TK}) E\big[v_1^T\big(\alpha^{(k,j')}\big) | \sigma^*, h^t_{\text{pub}}, m_{t^{(k,j)}} = \theta, \pi_{\min T^{(k,j)}} = \pi\big]$$
  $$= \bar{w}_1^{(k,j)}\big(h^t_{\text{pub}}\big) + \frac{1-\delta^{TK}}{\delta^{TK}} \tau^{(k,j)}(\theta).$$

Lemma S.5. *For any $\epsilon > 0$, there exists a $\bar{K} < \infty$ and $\bar{\delta} < 1$ such that if $K \geq \bar{K}$, $\delta \geq \bar{\delta}$, and $\delta^{TK} \leq 1 - \epsilon$, then in any (T,K)-recurrent mechanism in the game with an expert, σ^* is an* *optimal strategy. For any history h^t at the start of block (k, j) and any $j' > j$,*

$$\begin{aligned}&\sum_{t'\in T^{(k,j')}} (1-\delta)\delta^{t'} E\big[u_{t'}|\sigma^*, h^t\big] \\ &\quad \in B\Big(E\big[(1-\delta^T)\delta^{KTj'+kT} v^T\big(\alpha^{(k,j')}\big)|\sigma^*, h^t\big], (1-\delta^T)\delta^{KTj'+kT}\epsilon\Big).\end{aligned} \tag{S22}$$

Proof. For $\epsilon_1 > 0$, let $K_1 < \infty$ be such that for any prior, $\|\pi_{T(K_1-1)} - \pi\| < \epsilon_1$. Fix $\delta_1 < 1$ and $\epsilon_1 > 0$ such that if $\delta \geq \delta_1$ and $\|\nu - \pi\| < \epsilon_1$, then Lemmas S.2 and S.3 hold with bound $\epsilon > 0$.

Let h^t be a history in block (k, j) and let $\Sigma^{(k,j)}$ be the set of feasible strategies in the T-period mechanism implementing $\alpha^{(k,j)}$. Lemmas S.2 and S.3 imply that for all $j' > j$,

$$\left\| \sum_{t'\in T^{(k,j')}} \delta^{t'-\min\{T^{(k,j')}\}}(1-\delta)E\big[u_{t'}|\sigma^*, h^t\big] - (1-\delta^T)E\big[v^T\big(\alpha^{(k,j')}\big)|\sigma^*, h^t\big] \right\| \leq (1-\delta^T)\epsilon,$$

which in turn implies (S22).

It remains to show that σ^* is optimal for player 1. Consider a history h^t at the beginning of block (k, j) and define $d_1(m, \theta)$ as in (S12). Conditional on history h^t, θ_T is independent of σ. Thus, σ^* is optimal if (S13) holds.

If $\alpha^{(k,j)} \in \mathcal{A}^I \setminus \{\hat{\alpha}^{M,i}, \hat{\alpha}^{m,i}\}_{i=2}^N$, then (S13) holds by substituting $\hat{d}$ for d in the proof of Proposition 3. If $\alpha^{(k,j)} \in \{\hat{\alpha}^{M,i}, \hat{\alpha}^{m,i}\}_{i=2}^N$, note that (S15) holds. By Lemma S.3 it suffices to show that d_1 satisfies

$$d_1(m, \theta) - d_1(\hat{m}, \theta) \in B\big(\tau^{(k,j)}(m) - \tau^{(k,j)}(\hat{m}), \hat{d}(\delta, T)\big).$$

Following the same steps as Proposition 3, for any $\chi > 0$, ϵ_1 and ϵ_2 may be chosen sufficiently small that

$$\begin{aligned}d_1(m, \theta) - d_1(\hat{m}, \theta) \leq \sum_{j'=j+1}^{\infty} &\delta^{TK(j'-j)}\Big(E\big[v^T\big(\alpha^{(k,j')}\big)|\sigma^*, h^{\min\{T^{(k,j)}\}}, m_{t^{(k,j)}} = m\big] \\ &- E\big[v^T\big(\alpha^{(k,j')}\big)|\sigma^*, h^{\min\{T^{(k,j)}\}}, m_{t^{(k,j)}} = \hat{m}\big] + 2\chi\Big)\end{aligned}$$

with a similar bound from below. By property (ii) of Definition S.5, this bound may be written

$$\begin{aligned}d_1(m, \theta) - d_1(\hat{m}, \theta) &\leq \frac{\delta^{TK}}{1-\delta^{TK}}\left(\frac{1-\delta^{TK}}{\delta^{TK}}\big(\tau^{(k,j)}(m) - \tau^{(k,j)}(\hat{m})\big) + 2\chi\right) \\ &= \big(\tau^{(k,j)}(m) - \tau^{(k,j)}(\hat{m})\big) + \frac{2\delta^{TK}\chi}{1-\delta^{TK}},\end{aligned}$$

and similarly from below.

As long as $\delta^{TK} \leq 1-\epsilon$, $\chi > 0$ may be chosen sufficiently small (by choosing large $\bar{\delta} < 1$ and $\bar{K} < 1$) that $\frac{2\chi}{1-\delta^{TK}}$ is arbitrarily small. In particular, they can be chosen so that $d_1(m, \theta) \leq \hat{d}$. Therefore, σ^* is an optimal strategy by Lemma S.3. □

Definition S.6. A payoff $v \in \mathbb{R}^N$ is (T, ζ, W, δ)-decomposable in the game with an expert if there exists some $\alpha \in \hat{\mathcal{A}}^I$ and vectors $w(y,\theta) \in \mathbb{R}^N$ such that the following statements hold:

(i) If $\alpha \notin \{\hat{\alpha}^{M,i}, \hat{\alpha}^{m,i}\}_{i\in\{2,\dots,N\}}$, there exists $\bar{w} \in \mathbb{R}^N$ such that for all $\theta \in \Theta$, $E_y[w(y,\theta)|\alpha(\theta)] = \bar{w}$. If $\alpha \in \{\hat{\alpha}^{M,i}, \hat{\alpha}^{m,i}\}_{i\in\{2,\dots,N\}}$, there exists $\bar{w} \in \mathbb{R}^N$ such that $E_{\pi,y}[w(y,\theta)|\alpha(\theta)] = \bar{w}$.

(ii) The adding up constraint (6) holds.

(iii) Allocation rule α is (T, ζ, W, δ)-enforceable as follows:

(a) If $\alpha \notin \{\hat{\alpha}^{M,l}, \hat{\alpha}^{m,l}\}_{l=1}^N$, then for all i, m, and θ, (7) holds.

(b) If $\alpha \in \{\hat{\alpha}^{M,1}, \hat{\alpha}^{m,1}\}$, then for all m, θ, and $i \neq 1$, (7) holds. For $i = 1$, (7) holds if $m = \theta$.

(c) If $\alpha \in \{\hat{\alpha}^{M,l}, \hat{\alpha}^{m,l}\}_{l=2}^N$, then for all $\theta \in \Theta$, $E_y[w_1(y,\theta)|\alpha(\theta)] = \bar{w}_1 + (1-\delta) \times \tau^{\text{truth}}(\theta)$, where $\tau^{\text{truth}}(\theta)$ satisfies Assumption 4 for α. For all $i \neq l, m, \theta$, (7) holds. For $i = l$, (7) holds if $m = \theta$.

(iv) For any $\theta, \theta' \in \Theta$ such that $\alpha(\theta) = \alpha(\theta')$, for all $y \in Y$, $w(y,\theta) = w(y,\theta')$.

A set W is (T, ζ, δ)-decomposable if every $w \in W$ is (T, ζ, W, δ)-decomposable.

Lemma S.6. *Let $W \subseteq \mathbb{R}^N$ be a closed, convex, bounded set. Suppose there exists some $\hat{\delta} < 1$ such that for all $\delta \geq \hat{\delta}$, there exists a continuous function $\zeta(\delta) > 0$ such that W is $(T, \zeta(\delta), \delta)$ strictly self-decomposable. Then for all $\epsilon > 0$, there exists δ^* such that for all $\delta \geq \delta^*$ and $w \in W$, there exists an equilibrium of the infinite-horizon game with payoff $v \in B(w, \epsilon)$.*

Proof. Fix $\epsilon > 0$, let $W \subseteq \mathbb{R}^N$ be such a set, and consider $w \in W$. As in Proposition 3, there exists $\bar{\zeta} > 0$ such that if $\zeta < \bar{\zeta}$, there exists a $\bar{\delta} < 1$ such that for all $\delta \geq \bar{\delta}$, K can be chosen so that W is $(T, \bar{\zeta}, \delta^{TK})$-decomposable

Given δ and K, define $\hat{\delta} = \delta^{TK}$. Construct a mechanism as in the proof of Proposition 4. I claim this mechanism is (T, K)-recurrent in the game with an expert. Properties (i), (ii), and (iii) of Definition 8 follow immediately from the construction.

Consider properties (i) and (ii) of Definition S.5. Fix block (k, j). If $\alpha^{(k,j)} \in \{\hat{\alpha}^{M,i}, \hat{\alpha}^{m,i}\}_{i=2}^N$, then by definition of σ^*, player 1 reports truthfully in each period of $T^{(k,j)}$. Therefore,

$$E\big[w_1^{(k,j)}(y_{t^{(k,j)}}, m_{t^{(k,j)}})|\sigma^*, h^{\min\{T^{(k,j)}\}}, \pi_{\min T^{(k,j)}} = \pi\big]$$
$$= \bar{w}_1^{(k,j)} + \frac{1-\delta}{\delta} E_{\theta\sim\pi}\big[\tau(\theta)\big] = \bar{w}_1^{(k,j)}$$

because $E_{\theta\sim\pi}[\tau(\theta)]=0$. Applying (6) and noting that if $\pi_{\min T^{(k,j')}}=\pi$, then $\pi_{\min T^{(k,j'+1)}}=\pi$, I conclude

$$\delta^{TK}\bar{w}_1^{(k,j)}(h^t)=\sum_{j'=1}^{\infty}\delta^{TKj'}(1-\delta^{TK})E\big[v^T(\alpha^{(k,j+j')})|\sigma^*, h^{\min T^{(k,j+1)}}, \pi_{\min T^{(k,j+j')}}=\pi\big].$$

Therefore, property (ii) of Definition S.5 holds. If $\alpha^{(k,j)}\notin\{\hat{\alpha}^{M,i},\hat{\alpha}^{m,i}\}_{i=2}^N$, then a very similar argument shows that property (i) of Definition S.5 holds. So property (iv) holds and the constructed mechanism is (T,K)-recurrent.

Suppose that players can commit to actions as a function of messages. Lemma S.5 applies if $\delta<1$ and $K<\infty$ are sufficiently large. Hence, there exist $\delta_1<1$ and $K_1<\infty$ such that if $\delta\geq\delta_1$, $K\geq K_1$, and $\delta^{TK}<1-\epsilon$, then σ^* (from Definition S.5) is an optimal reporting strategy. If $\alpha^{(k,0)}\notin\{\hat{\alpha}^{M,i},\hat{\alpha}^{m,i}\}_{i=2}^N$, then $v\in B(w,2(1-\delta^{TK})+\delta^{TK}\zeta)$ as in Proposition 4. If $\alpha^{(k,0)}\in\{\hat{\alpha}^{M,i},\hat{\alpha}^{m,i}\}_{i=2}^N$, then consider a modified (T,K)-recurrent mechanism with $\alpha^{(k,0)}\in\hat{\mathcal{A}}^I\setminus\{\hat{\alpha}^{M,i},\hat{\alpha}^{m,i}\}_{i=2}^N$ and $\bar{w}^{(k,1)}=w$. The payoff v in this modified mechanism satisfies $\|v-w\|<2(1-\delta^{TK})+\delta^{TK}\zeta$ by construction. So $v\in B(w,2(1-\delta^{TK})+\delta^{TK}\zeta)$.

It remains to show that players have no incentive to deviate from the actions specified by the mechanism. If $\alpha^{(k,j)}\notin\{\hat{\alpha}^{M,i},\hat{\alpha}^{m,i}\}_{i=2}^N$, then the argument in Proposition 4 applies.[3] If $\alpha^{(k,j)}\in\{\hat{\alpha}^{M,i},\hat{\alpha}^{m,i}\}_{i=2}^N$, then player 1 may deviate in three ways. First, she could report truthfully but choose an incorrect action. Second, she could report falsely and play the correct action for her reported type. Third, she could report falsely and play the wrong action. The first and third types of deviation are not profitable because (7) holds for $i=1$ and all (m,θ). The second type of deviation is not profitable because truth-telling is an optimal strategy in the mechanism. Thus, player 1 has no profitable deviation.

Player $i\notin\{l,1\}$ likewise has no profitable deviation because (7) holds for all (m,θ). Player l believes that player 1 reports truthfully in each period and, hence, believes $m=\theta$ with probability 1. Therefore, player l has no profitable deviation because (7) holds for $m=\theta$.[4] This proves the result for $\delta^*>0$ and $\zeta>0$ such that $2(1-\delta^*)+\delta^*\zeta<\epsilon$. $\square$

Completing the Proof of Theorem 2. Statement (i) of Lemma S.1 goes through without change. For statements (ii) and (iii), the proof holds without change for $\alpha\notin\{\hat{\alpha}^{M,i},\hat{\alpha}^{m,i}\}_{i=2}^N$. If $\alpha\in\{\hat{\alpha}^{M,i},\hat{\alpha}^{m,i}\}_{i=2}^N$ and $\lambda\neq\pm(1,0,\ldots,0)$, then Lemma S.1 goes through if the targeted continuation payoff is $v_1^T(\alpha(\theta))+\frac{1-\delta}{\delta}\tau(\theta)$ for player 1 and $v_i^T(\alpha)$ for players $i\in\{2,\ldots,N\}$. If $\lambda=\pm(1,0,\ldots,0)$, then no bonus scheme makes α enforceable and

[3] Note that this argument applies regardless of players' prior over θ_t at the start of the period. Therefore, it holds regardless of whether players observe their own payoffs.

[4] Consider the setting in which players observe their own payoffs. Suppose that in period $t-1$, player l's payoff was inconsistent with player 1's reported type, even though player 1 was supposed to report truthfully with probability 1. So player l knows that player 1 has deviated. In that case, player 1 still has a strict incentive to report truthfully in the current period. Thus, define player l's beliefs following the deviation as any distribution with full support over types that are consistent with her observed payoff. Then incentives in the continuation game are identical to the case where l does not observe her payoff.

so $k^*(\alpha, \lambda, \zeta, \delta) = \mp\infty$. Note, however, that $k^*(\alpha, \lambda, \zeta, \delta)$ is the same as in Lemma S.1 if $\lambda = \pm(1, 0, \ldots, 0)$ and $\alpha \in \{\hat{\alpha}^{M,1}, \hat{\alpha}^{m,1}\}$.

For a unit normal λ, define $\hat{k}^T(\lambda) = \max_{\alpha \in \hat{\mathcal{A}}^I} \lambda \cdot v^T(\alpha)$. Let $\hat{H}^T(\lambda) = H^T(\lambda, \hat{k}^T(\lambda))$ and $\hat{Q}^T = \bigcap_\lambda \hat{H}^T(\lambda)$. If $\lambda = \pm(1, 0, \ldots, 0)$, then the proof of Proposition 5 uses $k^*(\alpha, \lambda, \zeta, \delta)$ only for $\alpha \in \{\hat{\alpha}^{M,1}, \hat{\alpha}^{m,1}\}$. Thus, the relevant k^* continues to be well defined and the proposition holds for $\hat{Q}^T$. But $\hat{V}^{T*} \subseteq \hat{Q}^T$ as in Theorem 1, so applying Lemma S.4 proves Theorem 2. □

D. Discussion of alternative mechanisms

This appendix discusses a natural alternative to (T, K)-recurrent mechanisms, and highlights why such a natural construction would not work in the proofs of Theorems 1 and 2. Consider the following alternative construction: separate the infinite-horizon game in blocks of T periods, each of which is followed by another block of T_A periods in which play is arbitrary. Fix $T_A > 0$ large. From the perspective of the last period in one block, the distribution after T_A periods is close to the invariant distribution (though not vanishingly close), so expected payoffs in that block are within some fixed $\epsilon > 0$ of invariant payoffs. In much of what follows, I ignore the T_A periods of arbitrary play and instead directly assume that expected payoffs in future blocks are no more than ϵ from invariant payoffs.

Let $B > 0$ be player i's myopic gain from a deviation in period t, and let D^j be the change in *invariant* payoffs in the block that is j blocks in the future from t (with $j \geq 1$) as a result of that deviation. For simplicity, assume expected payoffs in *each* future block are no more than ϵ away from invariant payoffs.[5] Then player i will not deviate if

$$(1-\delta)B + \sum_{j=1}^{\infty} \delta^{Tj}\big(1-\delta^T\big)\big(D^j + 2\epsilon\big) \leq 0$$

(2ϵ because both the on- and off-path expected payoffs could differ by ϵ from their respective invariant payoffs). Rewriting yields

$$(1-\delta)B + \sum_{j=1}^{\infty} \delta^{Tj}\big(1-\delta^T\big)D^j \leq -2\delta^T\epsilon.$$

To adapt the proof technique in FLM, it must be that $\sum_{j=1}^{\infty} \delta^{Tj}(1-\delta^T)D^j \to 0$ as $\delta \to 1$; otherwise, continuation payoffs would not be in the set W of payoffs to be approximated in equilibrium, since continuation invariant payoffs are drawn from a hyperplane that approaches the tangent hyperplane as $\delta \to 1$.

If $\epsilon = 0$ (as it does in FLM), then this limit poses no problems because $(1-\delta)B \to 0$. However, for any *fixed* $\epsilon > 0$, player i will prefer to deviate as $\delta \to 1$. Intuitively, player i cares much more about continuation payoffs than stage-game payoffs as $\delta \to 1$, so in particular cares about the possible gain in these continuation payoffs that arises due to

[5] This is a very loose bound because private information continues to deteriorate over time, but it is useful for illustrative purposes.

private information. In the limit, the potential gain from private information is larger than the incentives provided by continuation play, so player i cannot be deterred from deviating. To decrease ϵ, the construction must increase the number of inefficient periods T_A. But then $T_A \to \infty$ as $\delta \to 1$, so substantial inefficiencies might persist even in the limit. The (T, K)-recurrent mechanism avoids this problem because increasing K does not affect the efficiency of the resulting equilibrium.

References

Fudenberg, Drew, David Levine, and Eric Maskin (1994), "The folk theorem with imperfect public information." *Econometrica*, 62, 997–1039. [9]

Co-editor George J. Mailath handled this manuscript.

Manuscript received 11 August, 2014; final version accepted 1 October, 2016; available online 13 December, 2016.